REPORT DOCUMENTATION PAGE

AFRL-SR-BL-TR-01-

0002

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. F
		1 July 1996 - 31 December 1997

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
Updating Beliefs in Incompletely Specified Situations	F49620-96-1-0323

6. AUTHOR(S)
Prof. Halpern

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Cornell University Dept of Computer Science 4130 Upson Hall Ithaca, NY 14853	

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFOSR 801 N. Randolph Street, Room 732 Arlington, VA 22203-1977	F49620-96-1-0323

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION AVAILABILITY STATEMENT	12b. DISTRIBUTION CODE
Approved for Public Release.	

13. ABSTRACT (Maximum 200 words)

This project focused on a number of issues related to updating beliefs, both qualitative and quantitative, incompletely specified settings.

DTIC QUALITY INSPECTED 3

20010109 094

14. SUBJECT TERMS			15. NUMBER OF PAGES
			5
			16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT

Final Report: AFOSR Grant F49620-96-1-0323

1 Summary of research in 1996-97

This project focused on a number of issues related to updating beliefs, both qualitative and quantitative, in incompletely specified settings. In addition, work was initiated on the topic of explanation.

1.1 Qualitative Belief Revision

Nir Friedman and the PI provided a model for belief revision using plausibility measures During the past year, a great deal of time was spent polishing the results that appeared in conferences in previous years (and at the beginning of this year), both for inclusion in Nir Friedman's thesis and for submission for publication. In addition, new results were obtained. Here is a brief overview:

Belief revision, introduced by Alchourrón, Gärdenfors, and Makinson, focuses on how an agent revises his beliefs when he acquires new information. Katsuno and Mendelzon's *belief update*, on the other hand, focuses on how an agent should change his beliefs when he realizes that the world has changed. Both approaches attempt to capture the intuition that to accommodate the new belief the agent should make minimal changes to his beliefs. The difference between the two approaches comes out most clearly when we consider what happens when the agent observes something that is inconsistent with his previous beliefs. Revision treats the new observation as an indication that some of the previous beliefs are wrong and should be discarded. It tries to choose the most plausible beliefs that can accommodate the observation. Update, on the other hand, assumes that previous beliefs were correct and that the observation is an indication that a change occurred in the world. It tries to find the most plausible change that accounts for the observation.

Belief revision and belief update are just two points on a spectrum of possible belief change methods. There are situations where neither is appropriate. To investigate the problem of belief change more generally, it is useful to have a good formal model. Such a model is provided in [4]. We start with the model of knowledge in multi-agent systems introduced by Halpern and Fagin, and add to it (qualitative) plausibility to capture beliefs (where p is believed if its plausibility is greater than that of $\neg p$). Knowledge captures in a precise sense the non-defeasible information the agent has about the world he is in, while beliefs capture defeasible information. In [3], one point on the spectrum is examined, that can be viewed as combining the best features of revision and update. This approach assumes that the prior satisfies a (plausibilistic) Markovian assumption: that is, successive transitions are assumed to be independent, and the plausibility of a transition at time m depends only on the current global state, and not on what has happened up to time m.

Guided by these insights, we also re-examined the rationale behind standard approaches to belief revision, and found it wanting. In [2], we provided a critique of the literature on belief revision, stressing the importance of being explicit about the "ontology" or scenario underlying belief change.

1.2 Explanation

As probabilistic systems gain popularity and are coming into wider use, the need for a mechanism that explains the system's findings and recommendations becomes more critical. The system will also need a mechanism for ordering competing explanations. We examined two representative approaches to explanation in the literature—one due to Gärdenfors and one due to Pearl. An explanation for Gärdenfors is something that raises the probability of the *explanandum* (that which we are trying to explain). For Pearl, an explanation is the world description that has the highest

1

probability, given the explanandum. In [1], we show that both approaches suffer from significant problems. We proposed an approach to defining a notion of "better explanation" that combines some of the features of both together with more recent work by Pearl and others on causality.

1.3 Updating probabilistic information

Conditioning is the generally agreed-upon method for updating probability distributions when one learns that an event is certainly true. But it has been argued that we need other rules, in particular the rule of cross-entropy minimization, to handle updates that involve uncertain information [6]. We consider a well-known example of where cross-entropy might be used: van Fraassen's *Judy Benjamin* problem [7], which in essence asks how one might update given the value of a conditional probability. In [5], we argue that—contrary to the suggestions in the literature—it is possible to use simple conditionalization and thereby obtain answers that agree fully with intuition. This contrasts with proposals such as cross-entropy, which are easier to apply but can give unsatisfactory answers. This example again stresses the importance of providing an "ontology".

References

[1] U. Chajewska and J. Y. Halpern. Defining explanation in probabilistic systems. In *Proc. Thirteenth Conference on Uncertainty in Artificial Intelligence (UAI '97)*, pages 62–71, 1997.

[2] N. Friedman and J. Y. Halpern. Belief revision: A critique. In *Principles of Knowledge Representation and Reasoning: Proc. Fifth International Conference (KR '96)*, 1996.

[3] N. Friedman and J. Y. Halpern. A qualitative Markov assumption and its implications for belief change. In *Proc. Twelfth Conference on Uncertainty in Artificial Intelligence (UAI '96)*, pages 263–273, 1996.

[4] N. Friedman and J. Y. Halpern. Modeling belief in dynamic systems. part I: foundations. *Artificial Intelligence*, 1997. Accepted for publication. Also published as IBM Technical Report RJ 9965, 1995. A preliminary version appears in R. Fagin editor. *Theoretical Aspects of Reasoning about Knowledge: Proc. Fifth Conference*, 1994, pp. 44–64, under the title "A knowledge-based framework for belief change. Part I: foundations".

[5] A. J. Grove and J. Y. Halpern. Probability update: conditioning vs. cross-entropy. In *Proc. Thirteenth Conference on Uncertainty in Artificial Intelligence (UAI '97)*, pages 208–214, 1997.

[6] S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*. 22:76–86, 1951.

[7] B. C. van Fraassen. Symmetries of personal probability kinematics. In N. Rescher, editor, *Scientific Enquiry in Philsophical Perspective*, pages 183–223. University Press of America, Lanham. Maryland, 1987.

Joseph Y. Halpern's Activities: 1996-97

1 Honors

- Received 1997 Gödel Prize jointly with Yoram Moses for outstanding journal articles in the area of theoretical computer science over the past six years for paper "Knowledge and Common Knowledge in a Distributed Environment."

- Appointed editor-in-chief of *Journal of the ACM*

- Paper "Plausibility measures and default reasoning" commended for its excellence by the Committee on the "IGPL/FoLLI Prize for the Best Idea of the Year 1996".

2 University Activities

Co-director, Cognitive Studies Program

3 Professional Activities

- Fellow, American Association of Artificial Intelligence

- Editorial positions:

    - Editor-in-chief: *Journal of the ACM* (as of May, 1997)
    - Consulting Editor: *Chicago Journal of Computer Science*
    - On editorial board: *Artificial Intelligence Journal, Information and Computation, Journal of Logic and Computation*

- Member, ACM Publications Board (as of May, 1997)

- Chairman, ACM Preprint Repository Project (as of May, 1997)

- Program Committee Member: AAAI '97, UAI '97 2nd International Conference on Temporal Logic, Fourteenth National Conference on AI (AAAI '97)

- Conference Chair: 7th Conference on Theoretical Aspects of Rationality and Knowledge

- President of Board of Directors: Corporation for Theoretical Aspects of Reasoning About Knowledge

- Member, NSF Review Panel

4 Lectures

- Using multi-agent systems to represent uncertainty. Invited talk. AAAI '96 (Thirteenth National Conference on Artificial Intelligence), Portland, August, 1996.

- ——. IRCS, University of Pennsylvania, Philadelphia, PA, March, 1997

- ——. Invited talk, SCAI '97 (Sixth Scandinavian Conference on AI), Helsinki, August, 1997.

- Reasoning about knowledge in multi-agent systems. Economics Department, Cornell University, Ithaca, NY, September, 1996.

- —. Economics Department, University of Pennsylvania, Philadelphia, PA, March, 1997.

- —. C.S. Department, Indiana University, Bloomington, IN, May, 1997.

- —. C.S. Department, Washington University, St. Louis, MI, May, 1997.

- On ambiguities in the interpretation of game trees. SITE 96 Workshop on Game Theory, Stanford University, Palo Alto, CA, 1996.

- —. Economics Department, Cornell University, Ithaca, NY, May, 1997.

- From Statistics to Beliefs. Philosophy Dept., Carnegie-Mellon University, Pittsburgh, PA, May, 1997.

- Representation Independence in Probabilistic Reasoning. ISP, Pittsburgh University, Pittsburgh, PA, May, 1997.

- A counterexample to theorems of Cox and Fine, AAAI '96 (Thirteenth National Conference on Artificial Intelligence), Portland, Pittsburgh, PA, August, 1996

- Defining relative likelihood in partially-ordered preferential structures, Twelfth Conference on Uncertainty in AI, Portland, OR, August, 1996.

- Belief revision: a critique, 5th International Conference on Principles of Knowledge Representation and Reasoning (KR '96), Boston, MA, November, 1996.

5 Publications

1. From statistical knowledge bases to degrees of belief, *Artificial Intelligence* 87:1-2, October, 1996, pp. 75–143 (with F. Bacchus, A. J. Grove, and D. Koller).

2. Should knowledge entail belief?, *Journal of Philosophical Logic* 25:5, October, 1996, pp. 483-494.

3. A theory of knowledge and ignorance for many agents, *Journal of Logic and Computation* 7:1, January, 1997, pp. 79-108.

4. A critical reexamination of autoepistemic logic, default logic, and only knowing, *Computational Intelligence* 13:1, February, 1997, pp. 144-163.

5. On ambiguities in the interpretation of game trees, *Games and Economic Behavior* 20, 1997, pp. 66–96.

6. On the expected value of games with absentmindedness, *Games and Economic Behavior* 20, 1997, pp. 51–65 (with A. J. Grove).

7. Knowledge-based programs, *Distributed Computing* 10:4, 1997, pp. 199–225 (with R. Fagin, Y. Moses, and M. Y. Vardi).

8. Modeling belief in dynamic systems. Part I: Foundations, *Artificial Intelligence* 95:2, 1997, pp. 257–316 (with N. Friedman).

9. A counterexample to theorems of Cox and Fine, *AAAI-96* (Proceedings of the Thirteenth National Conference on Artificial Intelligence), August, 1996, pp. 1313–1319.

10. Plausibility measures and default reasoning, *AAAI-96* (Proceedings of the Thirteenth National Conference on Artificial Intelligence), August, 1996, pp. 1297–1304, to appear, *Journal of the ACM* (with N. Friedman)

11. First-order conditional logic revisited, *AAAI-96* (Proceedings of the Thirteenth National Conference on Artificial Intelligence), August, 1996, pp. 1305–1312 (with N. Friedman and D. Koller).

12. Irrelevance and conditioning in first-order probabilistic logic *AAAI-96* (Proceedings of the Thirteenth National Conference on Artificial Intelligence), August, 1996, pp. 569–576 (with D. Koller).

13. Using multi-agent systems to represent uncertainty, Summary of invited talk, *AAAI-96* (Proceedings of the Thirteenth National Conference on Artificial Intelligence), pp. 1329–1330, August, 1996.

14. Defining relative likelihood in partially-ordered preferential structures, *Proceedings of the Twelfth Conference on Uncertainty in AI*, August, 1996, pp. 299–306, and *Journal of AI Research.* , vol. 7, 1997, pp. 1–24.

15. A qualitative Markov assumption and its implications for belief change, *Proceedings of the Twelfth Conference on Uncertainty in AI*, August, 1996, pp. 263–273 (with N. Friedman).

16. Common knowledge: now you have it, now you don't, *Intelligent Systems: A Semiotics Perspective, Proceedings of the International Multidisciplinary Conference, Vol. 1*, October, 1996, pp. 177–183 (with R. Fagin, Y. Moses, M. Vardi).

17. Belief revision: a critique, *Proceedings of the 5th International Conference on Principles of Knowledge Representation and Reasoning (KR '96)*, November, 1996, pp. 421–431 (with N. Friedman).

18. Probability update: conditioning vs. cross-entropy, *Proceedings of the Thirteenth Conference on Uncertainty in AI*, 1997, pp. 208–214 (with A. Grove).

19. Defining explanation in probabilistic systems, *Proceedings of the Thirteenth Conference on Uncertainty in AI*, 1997, pp. 62–71 (with U. Chajewska).

6 Students

- Nir Friedman, Stanford (completed Ph.D., December, 1996)

- Urszula Chajewska, Stanford

- Francis Chu, Cornell